Get More With Less: Near Real-Time Image Clustering on Mobile Phones

Jorge Ortiz
IBM Research

Chien-Chin Huang
New York University

Supriyo Chakaborty
IBM Research

arXiv:1512.02972v1 [cs.CV] 9 Dec 2015

ABSTRACT

Machine learning algorithms, in conjunction with user data, hold the promise of revolutionizing the way we interact with our phones, and indeed their widespread adoption in the design of apps bear testimony to this promise. However, currently, the computationally expensive segments of the learning pipeline, such as feature extraction and model training, are offloaded to the cloud, resulting in an over-reliance on the network and under-utilization of computing resources available on mobile platforms. In this paper, we show that by combining the computing power *distributed over a number of phones*, judicious optimization choices, and contextual information it is possible to execute the end-to-end pipeline entirely on the phones at the edge of the network, efficiently. We also show that by harnessing the power of this combination, it is possible to execute a computationally expensive pipeline at near real-time.

To demonstrate our approach, we implement an end-to-end image-processing pipeline – that includes feature extraction, vocabulary learning, vectorization, and image clustering – on a set of mobile phones. Our results show a 75% improvement over the standard, full pipeline implementation running on the phones without modification – reducing the time to *one minute* under certain conditions. We believe that this result is a promising indication that fully distributed, infrastructure-less computing is possible on networks of mobile phones; enabling a new class of mobile applications that are less reliant on the cloud.

1. INTRODUCTION

Machine learning has revolutionized a broad range of fields and has the potential to change the way we interact with our mobile phones. The past few years has seen an increasing number of apps on phones that have exploited learning algorithms to provide context-aware services such as speaker recognition [15, 17], activity recognition [11], emotion detection [18], and several others [21, 23, 19]. However, several machine learning components such as feature extraction and model training are associated with a high computational and/or communication cost – consuming a large amount of energy and requiring high reliability – and are typically offloaded to the cloud. Cloud offloading [9, 8] provides several benefits to mobile applications. It not only provides a central location to gather data from *a large number* of phones [2], enhancing the quality of the results, but it also extends the battery life of mobile devices [8].

However, offloading of data also suffers from multiple drawbacks. First, data must be transmitted from the device to the cloud, potentially exposing a user's personal information, such as location traces or images. Second, depending on the status of the network connection the cloud-based service can become unavailable. Such network based outages are commonplace when many phones are co-located in a fixed geographic region, for example at stadium or an event in a park. The performance of a cellular connection depends on the number of users active in your "cell" – diminishing in quality as the service cell becomes congested [22]. Finally, there is disproportionate growth in data generation and additional network bandwidth. It is projected that global mobile data traffic will increase nearly tenfold between 2014 and 2019. While the 4G/LTE cellular network can increase its bandwidth by 20x, the projected demand will still exceed capacity in 2016. In response, many users offload computation through different channels to the internet. By 2016, more than half of all traffic from mobile devices (almost 14 exabytes) will be offloaded to the fixed network by means of Wi-Fi devices and femtocells each month [5].

We propose exploration within a different operating regime – one in which a cellular link is poor or intermittently available and users are incentivised to cooperate in order to send useful information to the cloud. For example, if users wish to send images from a popular protest, it would be useful for their phones to cooperatively choose which images to transmit and the order in which to transmit them, in order to make the best use of the intermittent cellular link to the cloud. Consider the approach when a good link is available. Phones could individually send their images to a server where a common *basis* space (for representing the images) is computed. The basis is used to then compress images into a bag-of-words (BoW) representation. Each BoW vector can be used to partition the images into k cluster using K-Means (Lloyd's algorithm) [13]. Although this is just one of many ways to cluster images, it is a commonly used pipeline for image clustering [7]. However, in the new

operating regime we consider, the network communication cost dominates. Therefore, we must execute this pipeline in the network itself, among the mobile phones. In this paper, we investigate the key question regarding the feasibility of the above proposal.

Cluster computing with mobile phones has been discussed in the past. The HPC community has envisioned scenarios where spare compute cycles can be used in a mobile cluster [4]. With power and communication being a concern, it is clear that not all algorithms would work well in this setting. However, with the increasing capabilities and compute power of phones, we believe it is time to revisit this question and examine it more thoroughly. For deeper analysis, we implement a distributed clustering application on phones, whereby the phones in the network collaborate to construct a representative subset of images to send back to a central repository. We show that a direct implementation of the cloud-based version yields very poor results but that several optimization makes the solution feasible. We achieve over a 4.3x improvement over the naive implementation with over *75% cluster similarity* with the full solution. We choose distributed image clustering as it is computationally expensive and data intensive providing benefits when the available bandwidth is low or unstable. The algorithm runs in stages with components whose runtime can be improved with contextual information, further providing opportunities for optimization and distribution. We also explore the design space by constructing a simple model based on the empirical measurements made on the implemented system and observe that dynamic pipeline shifting as a function of available bandwidth could provide improved performance as we move from a network-bound operating regime to a compute-bound one.

We believe that collaborative computation on mobile phones is under-explored and that there is an opportunity to leverage *high-density co-location* of mobile phones to support a new class of applications. We also believe that with the increasing power of mobile phones, now is the time to consider architectures and techniques that were previously unfeasible – even just a couple of years ago. Through our initial exploration, we show that by combining approximation algorithms, distributed computing, and contextual information, we can harness the power of machine learning at the edge of the network while reducing the computational overhead. In the rest of the paper we will describe our implementation, optimizations, and modeling results. We close with a discussion about their implication.

2. IMAGE CLUSTERING PIPELINE

Our image clustering pipeline is illustrated in Figure 1. It consists of four stages. The first stage runs scale-invariant feature transform (SIFT) [14] on each image to extract a set of features. The features from each image are treated as a single collection and passed to the next stage. The second stage runs K-Means on the collection and treats the *cluster heads* (centroids) as the representative 'words' or vocabulary of the dataset. The third stage labels the features for each image and generates a vector, where each row represents a specific word and the corresponding value is the number of occurrences of that word in the image. This is known as a bag-of-words presentation (BoW) [10]. Finally, each BoW vector is clustered again using K-Means. The clusters here represent the different 'types' of images in the data set. In our application, the image closest to the centroid for each cluster is transmitted to the cloud for storage. We consider the transmitted image set to be representative of the images in the network.

2.1 K-Means Algorithm

Given a set of n points $\{x_0, x_1, \ldots, x_{n-1}\}$ and k cluster heads (centroids) $\{c_0, c_1, \ldots, c_{k-1}\}$ the K-Means algorithm's goal is to minimize the following function:

$$J(X, C) = \sum_{j=1}^{k} \sum_{x_i \epsilon C_j} \parallel x_i - c_j \parallel^2 \quad (1)$$

This function is the square of the distance from each point within a cluster to the cluster's centroid, $c_j = \frac{1}{|C_j|} \sum_{x_i \epsilon C_j} x_i$. Minimizing this objective function is NP-hard [16]. However, the K-Means algorithm is a heuristic algorithm that is guaranteed to find a local minimum. The algorithm works as follows: given an initial value for k and a set of k random coordinates $\{c_0^0, c_1^0, \ldots, c_{k-1}^0\}$ the algorithm consists of the following steps:

1. **Assignment step** Given the current set of centroids, each point $\{x_0, x_1, \ldots, x_{n-1}\}$ is assigned to the corresponding cluster for which $x_i - c_j$ is smallest.
2. **Update step** Once each point is assigned to a cluster, recalculate the centroid coordinates by averaging all the points in a given cluster, $c_j = \frac{1}{|C_j|} \sum_{x_i \epsilon C_j} x_i$.

The first step is $O(ndk)$ and the second step is $O(n)$. The entire algorithm runs in $O(nkdi)$ where n is the number of points, k is the number of clusters, d is the dimensionality of x, and i is the number of iterations that the algorithm takes to converge. Convergence is defined with a stoppage criteria which is either 1) a fixed number of iterations, 2) no change in the cluster heads, or 3) a change in the cluster heads within some ϵ from one iteration to the next.

In our experiments, we run a distributed version of this algorithm. In distributed K-Means [12], a typical setup contains one master node that manages the iterations of the algorithms and determines when the process is complete. The master node broadcasts the starting centroids. Each node then runs the *assignment step* and the *update step* locally and reports their new centroid calculation back to the master. When the master receives the centroids from each of the nodes, it computes the average of each centroid and initiates another iteration. The process stops when the stoppage criteria is met.

2.2 Image Feature Extraction

For image feature extraction we use the SIFT technique that scans an image, looking for distinctive local features – present over a fixed location within the image itself. These typically indicate fluctuations in pixel values, like those found through corner-detection techniques. SIFT is robust to several kinds of transformations, such as scaling, rotation, affine, 3D perspective and various others. SIFT outputs a high-dimensional vector for each feature, known as a descriptor.

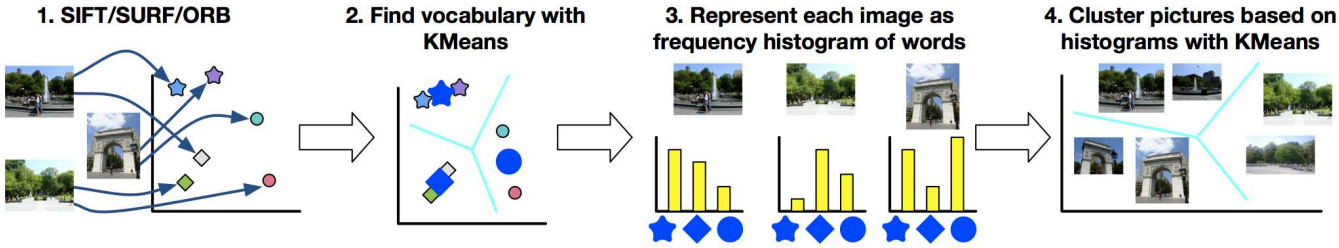
1. SIFT/SURF/ORB
2. Find vocabulary with KMeans
3. Represent each image as frequency histogram of words
4. Cluster pictures based on histograms with KMeans

Figure 1: Stages in the machine learning pipeline. First we run SIFT to extract features of the images, K-Means to construct the feature vocabulary for vectorization, we vectorize the images, and we run K-Means on those to cluster the images. Finally, the phones with images closest to each centroid uploads that image.

Because of the robustness of SIFT, it is typically used to look for similar features across images that contain overlapping scenes. Although feature extraction is robust to various kinds of transformations it still varies enough that similar features across pairs of images from different angles can vary slightly. To smooth out the variability we cluster them and use the centroids as the representative feature set. In our experiments, we tried several other feature-extracting techniques, such as SURF [6] and ORB [20] – with ORB giving us a measurable performance improvement. Indeed, either algorithm can be substituted for SIFT. However, in practice, both give worse clustering results. It is a fundamental tradeoff for execution time improvement.

2.3 Full Pipeline, In-Network Performance

We run the entire pipeline, illustrated in Figure 1, on a set of three mobile phones: Google Nexus 5, Motorola Nexus 6, and the Samsung Galaxy S4. The Nexus 5 has a Quad-core processor 2260 MHz Krait 400 with 32GB of storage. The Nexus 6 has a Quad-core 2.7 GHz Krait 450 also with 32GB of storage. The Samsung Galaxy S4 has a Quad-core 1900 MHz Krait 300 and 64 GB of storage. We collected a set of photos from each of the phones using different resolutions within a park in downtown Manhattan. We tried to collect a dataset that approximates one collected from a set of co-located phones.

We make a slight modification to a common processing pipeline used for classifying images [7]. The difference is in the last stage. Typically, a classifier such as Naive Bayes or SVM is used to classify the photos after they have been vectorized. We choose to cluster them instead, since our application aims to construct a representative set of photos. Our approach clusters the images and transmits the images closest to the centroids. The setting for k in both k-meeans (in the pipeline) was chosen experimentally. We visually assessed the quality of the output before choosing them. For the purposes of this study, we believe this is acceptable. Some amount of pre-processing must be done before running the pipeline in the network.

Figure 2 shows the execution time of each stage in the pipeline for a different number of images. 'SIFT Detection' denotes the SIFT computation stage which extracts the features from each image, 'Vocabulary K-Means' is the K-Means stage that is used to discover a representative set of features/words, 'BoW detection' is the vectorization step, and 'Picture K-Means' clusters the images using their vector

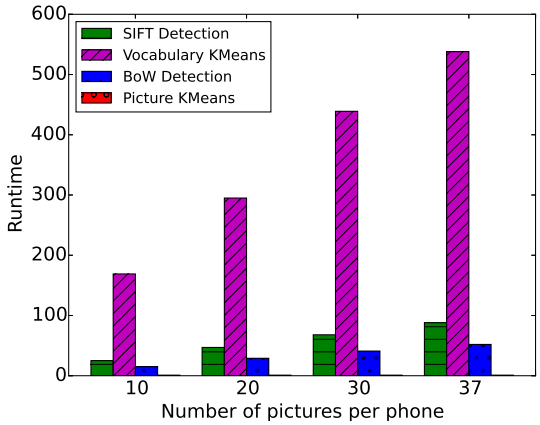

Figure 2: The figure shows that time is takes to execute each stage in the pipeline. Notice, the first K-Means stage – to construct the vocabulary from the extracted features – is the main bottleneck in overall performance.

representation. Note that for all stages, the 'Vocabulary K-Means' is by far the most expensive. With 37 images per phones, the entire pipeline takes an average 679 seconds (over 11 minutes) to complete; the second stage itself takes over eight minutes.

2.4 Optimizations

We implement a number of optimizations in order to decrease the execution time. The main challenge is maintaining clustering quality as we add more approximations. To measure the tradeoff, we compare our optimized pipeline cluster to the original one. Since the performance bottleneck is in the second stage of processing pipeline, we focus our attention on improving the execution time of K-Means. K-Means can be optimized a number of ways. Recall that the runtime of K-Means is O(ndki), where n is the number of points, d is their dimensionality, k is the number of clusters, and i is the number of iterations. We can optimize along any of these four parameters. In our work, we focus on the number of points and the number of iterations.

NDK vs SDK We implement our pipeline using the Android NDK [1]. The Android NDK allows you to write native code that runs on the device. It is typically used by applications that run computationally intensive jobs, pro-

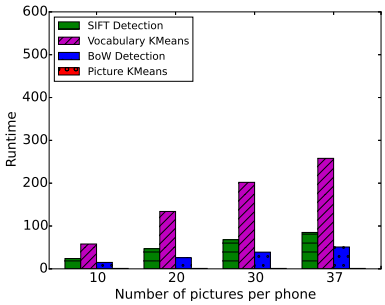

(a) Approx. K-Means with BoW.

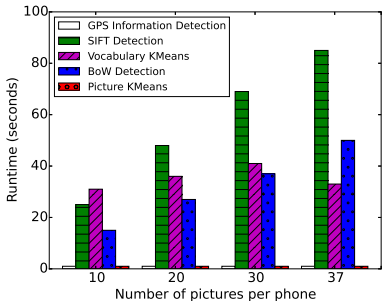

(b) EXIF with BoW.

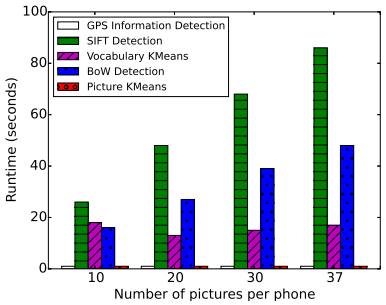

(c) Approx. K-Means with EXIF and BoW.

Figure 3: Performance for each pipeline stage for each combination of optimizations. The optimizations include context-based seeding with GPS and orientation data as well as a version of approximate K-Means that discards non-active points in the update step.

viding considerable performance gains. We used the NDK to write both our image processing code as well as our clustering code. Indeed, the performance gains are significant. Using 37 images per phone, on three phones, the completion time with the SDK is over two hours. The same pipeline, implemented with the NDK finishes in 679 seconds. Figure 2 shows the execution time of each pipeline stage with the NDK implementation. Notice that the second stage of the pipeline dominates the overall execution time.

Approximate K-Means There is a large body of work on approximate K-Means algorithms. Wang et al. [24] observe that in many datasets most of the points in the *assignment step* do not get re-assigned. The points that are most likely to get re-assigned are those that sit on the boundary between any cluster pair. They refer to these points as *active points*. Let $d(x_p, c_j)$ define the distance between point x_p and centroid c_j. Also, let $r = 1 - \frac{d(x_p,c_i)}{d(x_p,c_j)}$ define the relative ratio between the distance between the two closest centroid c_i and c_j. They show that active points – across a set of disparate data sets – have a distribution whereby over 90% of active points have an r-value < 0.15.

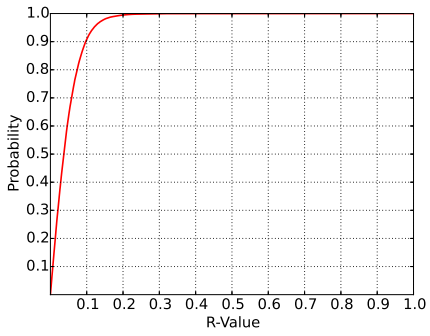

Figure 4: The r-value CDF for active points. Note, non-active points have an r-value of 1 and are no plotted in the figure. 90% of the active points in our dataset have an r-value 0.1 or less – similar to the distribution by Zeng et al. [24].

We calculate the r-value distribution in our dataset. Figure 4 shows a CDF of the r-value distribution for our *active points*. In our implementation, we identify the set of *active points* as the points that change cluster membership from the first iteration to the second one. In our dataset, 70% of points are considered *active points*; 90% of which have an r-value < 0.10. This distribution is similar to the one found in [24]. We discard inactive point in the remaining iterations to reduce the runtime of k-means further. With 37 photos, the completion time is 395 seconds, a 41% reduction from the original pipeline. Figure 3a shows the performance improvement for each stage of the pipeline when approximate K-Means is used. The average reduction is 45% for each experiment, as we vary the number of photos being processed.

Metadata Seeding Another useful approach to improve K-Means performance is to provide a good set of initial centroids. Our hypothesis is that location and orientation can be used to hint about which features may appear in the images. We recorded GPS and orientation in JPEG EXIF [3] metadata and ran K-Means on this data first. Then we used the centroid to seed to 'Vocabulary K-Means' pipeline stage. The EXIF K-Means is very inexpensive to run, as d is very small and i is also generally small (it converges quickly). We find that seeding the vocabulary-constructing K-Means stage (stage 2) significantly reduces the number of iterations and improves cluster quality. It also outperforms the original pipeline when random centroids are used. Figure 3b shows the results of metadata seeding. We see an overall average reduction of 70% in execution time from the original pipeline. We include the metadata seeding K-Means in our calculation.

Our original pipeline performs worse with a random set of centroids, on average, than with a hint from the metadata. In order to measure the improvement seen with the metadata, we compare the cluster overlap between a well-seeded full pipeline and one that is seeded with the metadata. To seed the original pipeline we use a prior run of that k-means step as our initial seed. We find that in over 100 runs, there is an average overlap of 74.7% with the well-seeded one. By comparison, the original pipeline with a random centroid seed overlaps with the well-seeded version by 69.3% on average. For our dataset, the 5% improvement is qualitatively observable.

Fully Optimized

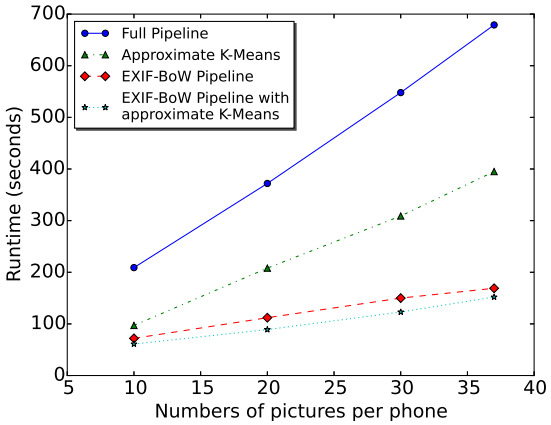

Figure 5: Performance measurements of entire pipelines on three phones. The execution time of the pipeline increases with the size of the input.

In our fully optimized pipeline, we combine both approximate K-Means and metadata seeding. Figure 3c shows the performance of each stage in the pipeline. The overall execution time for 37 photos is 152 seconds. The average performance reduction from the original is 75%. The largest performance improvement is seen by metadata seeding. Approximation adds another 5% reduction. With only 10 images, the overall execution time is 61 seconds. The entire pipeline can run in the mobile network and complete in *one minute*. We believe this makes it feasible and cheap enough to run as a background process for a family of mobile sensing applications, without interfering with phone interaction. Figure 5 shows a summary of each of the optimization and the original pipeline. We can see that the biggest performance improvement is seen with metadata seeding. We can also observe that the optimized pipeline scales better than the original. Finally, notice in Figure 3b and Figure 3c, the use of metadata reduces the execution time of K-Means to the point where *SIFT* become the bottleneck. Improvements to the feature extraction step could further drive down the execution time of the pipeline.

3. DISCUSSION

Today's edge computing frameworks, from IoT sensor devices to mobile phones are heavily reliant on the cloud. Most application architectures use the cloud for either offloading their computation or acting as a mediary between devices, altogether. For high-density co-location scenarios, such as a stadium event or protest, the connection to the cloud through the cellular network is unstable or often unavailable. More generally, the growth in the number of cloud-reliant devices through the spread of IoT and the explosive use of smart phones, strongly suggests that cloud-reliant architectures will not always be feasible, especially for applications where predictable quality of service and latency is important, cloud-based architectures cannot guarantee either [25]. Our initial results suggest that computationally expensive pipelines can be executed at the edge, entirely, and yield comparable performance through a combination of approximation techniques and contextual information. Moving forward, we believe there should be further exploration on architectures and techniques that achieve comparable performance at the edge to jobs executed in the cloud.

This does not however imply that the cloud should be entirely ignored. Instead, we believe that the cloud should be used more judiciously. Below, we examine the notion of dynamically shifting components of the pipeline between the cloud and the mobile network. We model our four-stage pipeline as a linear combination of the processing times on the cloud and in the mobile network and data transmission time. We further used representative workloads to estimate their average processing times on the cloud and the mobile network. Equation 2 shows the sum of the aggregate components that our model considers.

$$cost = t^{cloud}_{compute} + t^{mobile}_{compute} + t_{tx} \quad (2)$$

We calculated the cost across several pipeline configurations, shifting stages back and forth between the cloud and the mobile network as the available bandwidth changes.

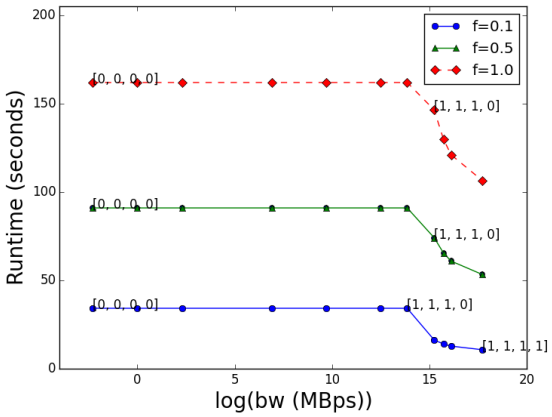

Figure 6: Model of pipeline shifting as a function of the available bandwidth. Note that as more bandwidth becomes available, shifting pipeline stages to the cloud yields better performance.

A value of 0 corresponds to placement of that stage in the mobile network and a 1 corresponds to placement in the cloud. For example, $[0, 0, 0, 0]$ means the entire pipeline runs in the mobile network while $[1, 1, 1, 1]$ means the entire pipeline runs in the cloud. Note, we only consider cases when the complete stage of a pipeline is shifted, i.e., pipeline shifting corresponds to a windowed placement shift for stages between the cloud and the mobile. Therefore we consider only five configurations ranging from entirely in the mobile network to entirely in the cloud. We also consider a different fraction of the data to send back by varying the fraction, f, of photos to transmit. Figure 6 shows the results of our simulation. We observe that shifting components of the pipeline could result in better performance. For future work we look to explore dynamic pipeline shifting in the context of an execution framework.

3.1 GPUs, Mobility, and Failure

There are several challenges that need to be addressed before we are able to fully realize the distributed computing

potential of the high density of mobile phones. It is important to study the effects of mobility and intermittent connectivity between these phones which will in turn effect the distribution of the tasks. We need to design robust protocols that will be able to handle these lower-level failures that are characteristics of the mobile networks but still being able to support higher-level synchronous tasks. In addition, there is also a constant increase in the computing power of these phones. GPUs could provide significant improvement in performance and energy consumption. However, there are implementation challenges and bottlenecks associated with copying data and objects from radio memory to CPU memory to GPU memorymaking the exact performance improvement unclear. We leave further exploration of this optimization for future work.

4. REFERENCES

[1] Android ndk. `https://developer.android.com/ndk/index.html`.

[2] The bright side of sitting in traffic: Crowdsourcing road congestion data. `https://googleblog.blogspot.com/2009/08/bright-side-of-sitting-in-traffic.html`.

[3] Exchangeable image file format for digital still cameras: Exif version 2.2. `http://www.exif.org/Exif2-2.PDF`.

[4] A cluster in your pocket:cell phone processors are getting more powerful. is a cell phone cluster possible? *Linux J.*, 2010(196), 2010.

[5] Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update 2014-2019. Technical report, Cisco, 01 2015.

[6] H. Bay, A. Ess, T. Tuytelaars, and L. Van Gool. Speeded-up robust features (surf). *Comput. Vis. Image Underst.*, 110(3):346–359, June 2008.

[7] G. Csurka, C. R. Dance, L. Fan, J. Willamowski, and C. Bray. Visual categorization with bags of keypoints. In *In Workshop on Statistical Learning in Computer Vision, ECCV*, pages 1–22, 2004.

[8] E. Cuervo, A. Balasubramanian, D. ki Cho, A. Wolman, S. Saroiu, R. Ch, and P. Bahl. Maui: Making smartphones last longer with code offload. In *In In Proceedings of ACM MobiSys*, 2010.

[9] S. Deng, L. Huang, J. Taheri, and A. Zomaya. Computation offloading for service workflow in mobile cloud computing. *Parallel and Distributed Systems, IEEE Transactions on*, PP(99):1–1, 2014.

[10] L. Fei-Fei and P. Perona. A bayesian hierarchical model for learning natural scene categories. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 2, pages 524–531. IEEE, 2005.

[11] N. D. Lane, Y. Xu, H. Lu, S. Hu, T. Choudhury, A. T. Campbell, and F. Zhao. Enabling large-scale human activity inference on smartphones using community similarity networks (csn). In *Proceedings of the 13th International Conference on Ubiquitous Computing*, UbiComp '11, pages 355–364, New York, NY, USA, 2011. ACM.

[12] Y. Liang, M. florina Balcan, and A. Kanchanapally. Distributed pca and k-means clustering.

[13] S. Lloyd. Least squares quantization in pcm. *Information Theory, IEEE Transactions on*, 28(2):129–137, Mar 1982.

[14] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision*, 60(2):91–110, Nov. 2004.

[15] H. Lu, A. J. B. Brush, B. Priyantha, A. K. Karlson, and J. Liu. Speakersense: Energy efficient unobtrusive speaker identification on mobile phones. In *Proceedings of the 9th International Conference on Pervasive Computing*, Pervasive'11, pages 188–205, Berlin, Heidelberg, 2011. Springer-Verlag.

[16] M. Mahajan, P. Nimbhorkar, and K. Varadarajan. The planar k-means problem is np-hard. In *Proceedings of the 3rd International Workshop on Algorithms and Computation*, WALCOM '09, pages 274–285, Berlin, Heidelberg, 2009. Springer-Verlag.

[17] E. Miluzzo, C. T. Cornelius, A. Ramaswamy, T. Choudhury, Z. Liu, and A. T. Campbell. Darwin phones: The evolution of sensing and inference on mobile phones. In *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services*, MobiSys '10, pages 5–20, New York, NY, USA, 2010. ACM.

[18] K. K. Rachuri, M. Musolesi, C. Mascolo, P. J. Rentfrow, C. Longworth, and A. Aucinas. Emotionsense: A mobile phones based adaptive platform for experimental social psychology research. In *Proceedings of the 12th ACM International Conference on Ubiquitous Computing*, UbiComp '10, pages 281–290, New York, NY, USA, 2010. ACM.

[19] S. Reddy, M. Mun, J. Burke, D. Estrin, M. Hansen, and M. Srivastava. Using mobile phones to determine transportation modes. *ACM Trans. Sen. Netw.*, 6(2):13:1–13:27, Mar. 2010.

[20] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski. Orb: An efficient alternative to sift or surf. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 2564–2571, Nov 2011.

[21] J. P. Rula, B. Jun, and F. Bustamante. Mobile ad(d): Estimating mobile app session times for better ads. In *Proceedings of the 16th International Workshop on Mobile Computing Systems and Applications*, HotMobile '15, pages 123–128, New York, NY, USA, 2015. ACM.

[22] R. Schoenen, A. Otyakmaz, and Z. Xu. Resource allocation and scheduling in fdd multihop cellular systems. In *Communications Workshops, 2009. ICC Workshops 2009. IEEE International Conference on*, pages 1–6, June 2009.

[23] Y. Wang, J. Yang, Y. Chen, H. Liu, M. Gruteser, and R. P. Martin. Tracking human queues using single-point signal monitoring. In *Proceedings of the 12th Annual International Conference on Mobile Systems, Applications, and Services*, MobiSys '14, pages 42–54, New York, NY, USA, 2014. ACM.

[24] G. Zeng. Fast approximate k-means via cluster closures. In *Proceedings of the 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, CVPR '12, pages 3037–3044, Washington, DC, USA, 2012. IEEE Computer Society.

[25] B. Zhang, N. Mor, J. Kolb, D. S. Chan, K. Lutz, E. Allman, J. Wawrzynek, E. Lee, and J. Kubiatowicz. The cloud is not enough: Saving iot from the cloud. In *7th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 15)*, Santa Clara, CA, July 2015. USENIX Association.